



Defence Research and
Development Canada Recherche et développement
pour la défense Canada



Training effects on the use of simple heuristics in threat assessment

David J. Bryant

Defence R&D Canada – Toronto

Technical Report

DRDC Toronto TR 2005-229

August 2005

Canada

Training effects on the use of simple heuristics in threat assessment


David J. Bryant

Defence R&D Canada – Toronto

Technical Report

DRDC Toronto TR 2005-229


August 2005

Author

David J. Bryant, Ph.D.

Approved by

Carol McCann

Head, Command Effectiveness and Behaviour Section

Approved for release by

K.M. Sutton

Chair, Document Review and Library Committee

Abstract

Previous research has examined the use of fast and frugal heuristics for threat classification with probabilistic cues. In all previous studies subjects learned the underlying relationships of cues to friend/foe classification through trial-and-error learning. Explicit training, however, is theoretically interesting because it potentially involves the integration of deliberate cognitive learning of the task's underlying cue structure with implicit, procedural learning of specific cue-criterion probabilities. In the experiment, subjects learned to classify contacts in a simulated naval warfare environment and then were tested on sets of contacts that were designed to contrast predictions of several heuristics, including the Take-the-Best-for-Classification (TTB-C) and Pros Rule developed specifically for the threat classification task, as well as a Bayesian strategy based on computation of the conditional probabilities of friend or foe classification given the particular pattern of cues. The interpretability of the results was limited by the large proportion of subjects who exhibited uninterpretable patterns of responding. In contrast to previous experiments, very few subjects employed TTB-C, although more did use the less frugal Pros Rule. The experiment did yield the novel finding that some subjects can, given explicit training, employ a Bayesian strategy for this task.

Résumé

Des recherches antérieures ont porté sur l'utilisation de l'heuristique simple et rapide pour la classification des menaces à l'aide d'indices probabilistes. Dans toutes les études précédentes, les sujets ont appris à discerner les relations sous-jacentes entre les indices servant à la classification ami/ennemi dans le cadre d'une approche par essais et erreurs. Une formation explicite est cependant intéressante sur le plan théorique parce qu'elle peut permettre d'intégrer l'apprentissage cognitif délibéré de la structure des indices sous-tendant la tâche et l'apprentissage procédural implicite des probabilités de certains indices-critères. Dans l'expérience, les sujets ont appris à classer les contacts dans un environnement simulé de guerre navale et ils ont fait l'objet d'un test sur des séries de contacts conçues pour mettre en parallèle les prédictions de plusieurs approches heuristiques, dont celle de « ne garder que le meilleur en vue de la classification » (TTB-C) et la « règle des pour » mise au point expressément pour la tâche de classification des menaces, de même qu'une stratégie bayésienne basée sur le calcul des probabilités conditionnelles de la classification ami/ennemi, compte tenu de la configuration particulières des indices. L'interprétabilité des résultats était limitée en raison de la forte proportion de sujets qui présentaient des modes de réponse non interprétables. Contrairement aux expériences antérieures, très rares étaient les sujets qui utilisaient la TTB-C, mais un plus grand nombre ont appliqué la règle des pour, qui est moins simple. L'expérience a fait ressortir l'élément nouveau suivant, à savoir que certains sujets peuvent, s'ils reçoivent une formation explicite, adopter une stratégie bayésienne pour cette tâche.

This page intentionally left blank.

Executive summary

This report presents the results of an experiment as part of an ongoing line of research investigating the cognitive mechanisms of threat classification. Given the nature of the sensors available to collect data and the composition of objects in an area of operation, threat assessment entails the classification of objects on the basis of cues that are probabilistic, generally not offering unequivocal indication of an object's status as friend, foe, or neutral.

Previous research has examined the extent to which people employ analytic or heuristic methods for threat classification. Early analytic approaches to explaining decision making were based on normative theories of probability and logic, whereas intuitive methods encompass a range of informal heuristics. Fast and frugal heuristics, for example, are simple procedures that are adaptive within the limits of time and knowledge imposed by the situation and the computational power and the decision maker. Studies have shown fast and frugal heuristics to be accurate and efficient solutions to certain judgment tasks. Studies have also discovered evidence that fast and frugal heuristics can serve as good models for actual human choice behaviour.

To apply the fast and frugal heuristic approach to threat assessment, Bryant [1] [2] devised a variant of TTB, called Take-the-Best-for-Classification heuristic (TTB-C). In previous experiments, it was found that a proportion of subjects employed TTB-C but other subjects employed a somewhat less frugal Unweighted Pros Rule, in which all cues are aggregated to determine which option has the greater support. No subjects were observed to employ a computationally intensive Bayesian strategy.

In all previous studies, subjects learned the underlying relationships of cues to friend/foe classification through trial-and-error learning. Although this type of training is common in probabilistic cue learning studies, people are often exposed to explicit training in real-world tasks. Explicit training is theoretically interesting because it potentially involves the integration of deliberate cognitive learning of the task's underlying cue structure with implicit, procedural learning of specific cue-criterion probabilities. This experiment investigates the impact of explicit information about cue information on the use of decision heuristics in the simulated air threat assessment task. Specifically, the experiment contrasts *pattern-based* training, in which typical patterns of cues are used to distinguish friend from foe, and *cue-based* training, in which detailed diagnostic information about each individual cue is presented.

In the experiment, participants learned to classify contacts in a simulated naval warfare environment and then were tested on the task under either a low or high time pressure. The test sets of contacts were designed to contrast predictions of several heuristics, including the Take-the-Best-for-Classification (TTB-C) and Pros Rule developed specifically for the threat classification task, as well as a Bayesian strategy based on computation of the conditional probabilities of friend or foe classification given the particular pattern of cues.

The results demonstrated that explicit training is suitable for learning to perform the assessment task. Subjects were able to perform to a near optimal level of accuracy immediately without extensive practice with both cue- and pattern-based training. In contrast to previous experiments, however, a small proportion of subjects were observed to use a

Bayesian strategy. Despite the previous findings that many subjects employed TTB-C in this task after feedback training, only three instances of subject data consistent with TTB-C were observed in this experiment. Other subjects employed either the Weighted or Unweighted Pros Rules. The results of the experiment were, unfortunately, limited by the high proportion of subjects who exhibited uninterpretable patterns of responding.

Although it was hypothesized that pattern-based training might lead to better cue integration and greater use of a Bayesian strategy than cue-based training, the results offer no evidence to support this hypothesis. Thus, the results of the current experiment do not offer a clear picture of how explicit training affects probabilistic cue learning and decision strategy use by subjects in a simulated threat assessment task. Nor do the results indicate how different kinds of explicit information might differ in their effects. Further experimentation, however, is warranted as the experiment did yield a novel finding that some subjects can, given explicit training, employ a Bayesian strategy for this task.

Bryant, D.J. 2005. Instructional effects on the use of simple heuristics in threat assessment. DRDC Toronto TR 2005-229 Defence R&D Canada - Toronto.

Sommaire

Le présent rapport rend compte des résultats d'une expérience effectuée dans le cadre d'une série continue de recherches portant sur les mécanismes cognitifs de la classification des menaces. Étant donné la nature des capteurs disponibles pour recueillir des données et la composition des objets dans une zone d'opération, l'évaluation des menaces comporte une classification des objets à partir d'indices probabilistes qui n'offrent pas généralement d'indication non équivoque du statut d'un objet comme ami, ennemi ou neutre.

Des recherches antérieures ont tenté de déterminer la mesure dans laquelle les personnes utilisent des méthodes analytiques ou heuristiques pour classer les menaces. Les premières approches analytiques pour expliquer le processus décisionnel se basaient sur des théories normatives de la probabilité et de la logique, alors que les méthodes intuitives englobent un éventail d'approches heuristiques informelles. L'heuristique simple et rapide, par exemple, regroupe des procédures simples qui peuvent être adaptées à l'intérieur des limites de temps et des connaissances imposées par la situation, la capacité de traitement et le décideur. Des études ont montré que l'heuristique simple et rapide apportait des solutions exactes et efficaces à certaines tâches faisant appel au jugement. Des études ont également révélé que l'heuristique simple et rapide peut offrir des modèles intéressants pour le comportement décisionnel humain.

Pour appliquer l'approche heuristique simple et rapide à l'évaluation des menaces, Bryant [1] [2] a mis au point une variante de la classification TTB, une heuristique consistant à « ne garder que le meilleur en vue de la classification » (Take-the-Best-for-Classification (TTB-C)). Dans des expériences précédentes, on a constaté qu'une proportion de sujets utilisaient la TTB-C mais que d'autres sujets avaient recours à une « règle des pour » non pondérée un peu moins simple, où tous les indices sont combinés pour déterminer quelle option reçoit le plus d'appui. On a observé qu'aucun sujet n'employait une stratégie bayésienne requérant beaucoup d'effort computationnel.

Dans toutes les études précédentes, les sujets ont appris à discerner les relations sous-jacentes entre les indices pour la classification ami/ennemi dans le cadre d'une approche par essais et erreurs. Bien que ce type de formation soit fréquent dans les études d'apprentissage des indices probabilistes, les personnes sont souvent soumises à une formation explicite sur les tâches dans le monde réel. La formation explicite est intéressante sur le plan théorique parce qu'elle peut permettre d'intégrer l'apprentissage cognitif délibéré de la structure des indices sous-tendant la tâche et l'apprentissage procédural implicite des probabilités de certains indices-critères. La présente expérience examine l'impact d'une formation explicite relative aux indices sur l'utilisation de l'heuristique décisionnelle dans le cadre d'une tâche simulée d'évaluation des menaces aériennes. Plus précisément, l'expérience met en parallèle la formation *basée sur des modèles*, où des configurations typiques d'indices sont utilisées pour distinguer un ami d'un ennemi, et la formation *basée sur des indices*, où une information diagnostique détaillée sur chaque indice individuel est présentée.

Dans l'expérience, les participants ont appris à classer les contacts dans un environnement simulé de guerre navale et ont ensuite été soumis à un test visant à évaluer la tâche avec une

contrainte temporelle faible ou forte. Les séries de contacts dans le test ont été conçues de façon à mettre en parallèle les prédictions de plusieurs approches heuristiques, dont celle de « ne garder que le meilleur en vue de la classification » (TTB-C) et la « règle des pour » mise au point expressément pour la tâche de classification des menaces, de même qu'une stratégie bayésienne basée sur le calcul des probabilités conditionnelles de la classification ami/ennemi, compte tenu de la configuration particulière des indices.

Les résultats ont montré qu'une formation explicite convient à l'apprentissage de la tâche d'évaluation. Avec la formation basée sur les indices et les modèles, les sujets ont pu obtenir un degré d'exactitude quasi optimal tout de suite sans longue pratique. Contrairement aux expériences précédentes, cependant, une faible proportion de sujets ont employé une stratégie bayésienne. Malgré les conclusions antérieures suivant lesquelles de nombreux sujets avaient recours à la TTB-C dans cette tâche après une formation utilisant le principe de la rétroaction, on a observé dans cette expérience seulement trois cas où les données sur les sujets correspondaient à la TTB-C. D'autres sujets ont utilisé soit la règle des pour pondérée ou non pondérée. Les résultats de l'expérience étaient malheureusement limités à cause de la forte proportion de sujets qui ont présenté des modes de réponse non interprétables.

Bien que nous ayons avancé l'hypothèse que la formation basée sur des modèles puisse contribuer à améliorer l'intégration des indices et à accroître l'utilisation d'une stratégie bayésienne par rapport à une formation basée sur les indices, les résultats ne viennent pas corroborer cette hypothèse. Les résultats de l'expérience actuelle ne donnent pas une image claire de la façon dont la formation explicite influe sur l'apprentissage des indices probabilistes et l'emploi de stratégies décisionnelles par les sujets dans une tâche simulée d'évaluation des menaces. Les résultats n'expliquent pas non plus comment différents types d'information explicite peuvent avoir des effets différents. D'autres expériences s'imposent toutefois car la présente expérience a fait ressortir un élément nouveau, à savoir que certains sujets peuvent, s'ils reçoivent une formation explicite, adopter une stratégie bayésienne pour cette tâche.

Bryant, D.J. 2005. Instructional effects on the use of simple heuristics in threat assessment. DRDC Toronto TR 2005-229 Defence R&D Canada - Toronto.

Table of contents

Abstract.....	i
Résumé	i
Executive summary	iii
Sommaire.....	v
Table of contents	vii
List of figures	ix
List of tables	ix
Acknowledgements	x
Introduction	1
Background.....	1
The Take-the-Best-for-Classification heuristic	2
Compensatory heuristics	3
Explicit training.....	5
Purpose of study	5
Method.....	7
Subjects	7
Materials	7
Design.....	7
Procedure	10
Results	11
Training session.....	11
Classification strategy	12
Response time.....	14
Cue selections.....	15

Discussion.....	17
References	19
Annex A.....	21
List of symbols/abbreviations/acronyms/initialisms	28
Distribution list.....	29

List of figures

Figure 1: The Take-the-Best-for-Classification (TTB-C) Heuristic	3
Figure 2: Weighted Version of the Pros Rule for Classification	4
Figure 3: Classification Accuracy by Block in the Training Session	11
Figure 4: Average Number of Cues Selected by Sequence	16

List of tables

Table 1: Relative Frequencies of Cue Values for Friend and Foe Contacts	8
Table 2: Predicted Responses to Contact Types by Hypothesized Strategies	9
Table 3: Predicted Accuracy Levels by Contact Type	9
Table 4: Response Accuracies by Contact Type and Decision Strategy Classification	13
Table 5: Number of Subjects Classified as Using Hypothesized Decision Strategies	14
Table 6: Average Number of Cue Selections by Decision Strategy	15

Acknowledgements

The experiment reported here was completed with the invaluable assistance of Andrea Hawton and Elaine Maceda.

Introduction

Background

This report presents the results of an experiment as part of an ongoing line of research investigating the cognitive mechanisms of threat classification. The threat classification task is a critical part of situation assessment and force protection. It entails the classification of objects on the basis of sensor data that provide cues as to the type and identity of each object. Given the nature of the sensors available and the composition of objects in an area of operation, these cues are probabilistic, generally not offering unequivocal indication of an object's status as friend, foe, or neutral. For this reason, previous experiments [1] [2] have examined how subjects learn to classify simulated air contacts in a low-fidelity air-threat environment.

In large part, the purpose of previous research was to determine whether people tend to employ analytic or heuristic methods for threat classification. The distinction between these methods is important in understanding how humans within the Command and Control (C2) system take in and evaluate uncertain information [3]. Early analytic approaches to explaining decision making were based on the premise that human decision making can be modeled in terms of formal processes predicted by normative theories of probability and logic [4]. Numerous sequential and distributed procedures for comparing alternatives are known, most of which can be computationally modeled by production systems operating on a representation of the problem space. Many, for example, are based on Bayesian theorems [5].

Analytic models, however, often fail to adequately describe peoples' decision making behaviours when performing real-world tasks. Gigerenzer, Todd, and the Adaptive Behavior and Cognition Group [6] have proposed that *fast and frugal heuristics* can serve as models for human judgment. The fast and frugal heuristic approach is based on a conceptualization of rationality in which behaviour is evaluated in terms of its adaptiveness within the limits of time and knowledge imposed by the situation and the computational power and the decision maker [7] [8]. Todd and Gigerenzer [9] define this concept of *ecological rationality* as "adaptive behavior resulting from the fit between the mind's mechanism and the structure of the environment in which it operates."

Studies have shown fast and frugal heuristics to be accurate and efficient solutions to certain judgment tasks. In addition to achieving levels of accuracy comparable to computationally intensive analytic procedures, fast and frugal heuristics consistently exhibit a clear advantage over linear procedures in terms of frugality, consulting, on average, fewer cues and performing fewer computations.

Studies have also discovered evidence that fast and frugal heuristics can serve as good models for actual human choice behaviour. For example, the Take-the-Best (TTB) heuristic, which chooses from two alternatives on the basis of the single most valid available cue, can be an effective strategy when the task structure affords several highly predictive cues. Several studies have examined subjects' choice behaviour in tasks in which subjects are required to

use probabilistically predictive cues to select an alternative (e.g., [10] [11] [12] [13]). Typically, these studies report that a proportion of subjects can be classified as employing a simple heuristic, such as TTB, to make choices, although a significant, often majority, proportion of subjects seem to use more complex, compensatory procedures. The propensity of subjects to employ a heuristic such as TTB is affected by a range of factors, such as costs imposed on obtaining cues. Broder [10] found 40% of subjects used TTB when the cost of cue information was relatively low and 60% when the cost was relatively high. Newell and Shanks [13] have also replicated the finding that a greater proportion of subjects employ TTB when cue costs are relatively higher.

The Take-the-Best-for-Classification heuristic

To apply the fast and frugal heuristic approach to threat assessment, Bryant [1] [2] devised heuristics specifically for the threat classification task. The key elements of fast and frugal heuristics are the frugal search (using the minimum of data) and simple decision rule (avoiding complex computations), both of which can be applied to classification as well as choice tasks. Threat classification can be performed by a heuristic that adheres to the principle of frugal search. In this case, the minimum information needed is one cue that can indicate the threat class to which a contact likely belongs. The task can also be performed with a simple decision rule of selecting the threat class to which the value of that cue is most strongly associated. The stopping rule for threat classification is built into the decision rule; i.e. search is terminated when a cue is located that can be used to make a decision.

A variant of TTB, called Take-the-Best-for-Classification heuristic (TTB-C), was devised to perform the threat classification task.¹ It is based on the premise that the single best cue can be used to make accurate threat classification judgments in a task environment in which that cue is highly predictive. Thus, TTB-C is not intended to be universally applicable but a fast and frugal alternative when one or more cues point to the appropriate threat classification at an acceptable rate.² Unlike TTB, which chooses between two objects along a single dimension, TTB-C places a single object into one of two categories along the threat dimension, as is illustrated in Figure 1.

TTB-C assumes that there exist one or more cues that have some predictive association to the threat class of contacts and that all, or some subset, of these cues can be inspected by the decision maker. Moreover, the decision maker must have acquired, through experience or training, knowledge of the relative validities of these cues. In all previous experiments [1] [2], people have been able to learn to classify contacts on the basis of probabilistic cues to a high degree of accuracy. Generally, participants in these experiments reached levels of accuracy only slightly less than the optimum level predicted by the normative Bayesian model.

¹ TTB-C is also derivable from the Lexicographic heuristic for two-alternative choice, which is a generalization of Take-the-Best [6; p. 143].

² What is an acceptable rate must be determined through the balancing of accuracy demands and resource limitations in terms of time and available information.

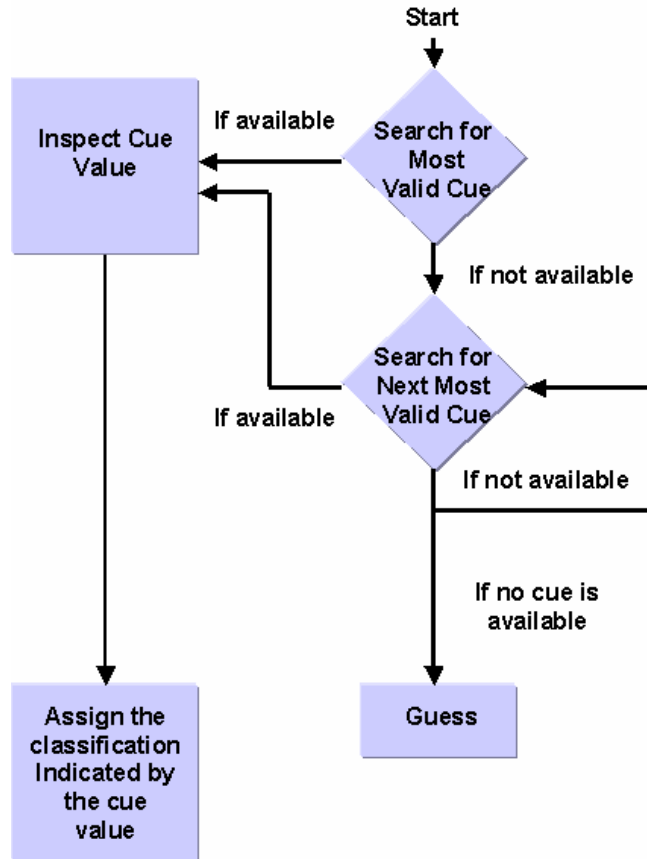


Figure 1. *The Take-the-Best-for-Classification (TTB-C) Heuristic*

Given an as-yet-unclassified contact, the heuristic begins by searching for the single most valid cue to serve as the basis for classification. In the experiment described in this report, all cues associated with contacts will be available, so the most valid cue should always be inspected. If the experimental procedure made some cues unavailable for certain contacts, the decision maker would have to determine whether the most valid cue was available and, if it was not, then search for the next most valid cue. The decision rule is equally simple; when the most valid available cue is located, the decision maker assesses which threat class has the greater probability of being true given the value of that cue and makes that threat class the output of the heuristic. The heuristic will be used to make the simplified two-category choice (friend or foe) of the experimental task but could apply to threat classification with the traditional set of threat classes (hostile, potential threat, neutral, and friend). With the contact classified, the heuristic terminates. Should no valid cue be found, the decision maker must guess.

Compensatory heuristics

Bryant [2] found that some subjects appeared to use TTB-C to make threat classification judgments but also that other subjects appeared to use a different, compensatory heuristic. A compensatory heuristic is one that employs all available cues, making it less frugal than TTB-C which is non-compensatory. A compensatory heuristic, however, can employ a simple

decision rule and serve as a fairly simple method to perform threat assessment – less complex, for example, than an analytic procedure that involves computing the Bayesian conditional probabilities of friend and foe classifications given the cues available.

The compensatory heuristic that some subjects employed was a variant of Dawes' Rule for choice. Dawes' Rule is a procedure by which a decision maker calculates the sum of cue values for each alternative and selects the alternative with the highest score [14]. A related heuristic is Franklin's Rule by which a decision maker calculates the sum of cue values weighted by the corresponding cue validities for each alternative and selects the alternative with the highest score. Because these rules involve accumulation of evidence for or against one alternative, they can be termed Pros Rules; i.e. according to the rule, the decision maker counts the number of affirmative cues, or pros, for each alternative and selects the one with the most pros associated with it. For the sake of clarity, Franklin's and Dawes' Rules will be referred to as the Weighted Pros and Unweighted Pros rules, respectively, throughout this report.

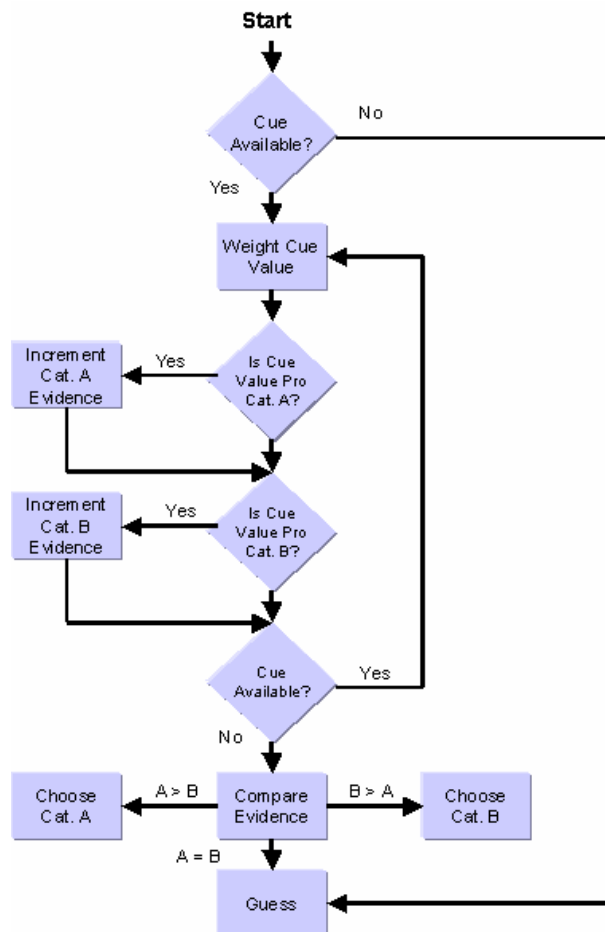


Figure 2: Weighted Version of the Pros Rule for Classification

To perform the threat classification task, the Pros Rule examines each cue value and assigns evidence toward either friend or foe classification, depending on the associations of cue

values to threat class. A running sum is maintained and, after all available cues have been inspected, used to place the contact in the friend or foe category. Figure 2 contains an illustration of Weighted Pros Rule, which weights cues by their validity, adapted for threat classification. A classification version of the Unweighted Pros Rule is performed just as illustrated in Figure 2 but without the weighting step following the selection of a cue. These rules use more information than TTB-C but are more generally useful because their accuracy is not limited to cases where a single cue is highly predictive.

Explicit training

Although feedback training, in which subjects learn the underlying cue structure without any explicit guidance, is common in probabilistic cue learning studies (e.g., [15]), people are often exposed to explicit training in real-world tasks. Explicit training is theoretically interesting because it potentially involves the integration of deliberate cognitive learning of the task's underlying cue structure with implicit, procedural learning of specific cue-criterion probabilities.

Juslin et al. [15] reported that integration of cues in a cue learning task is more likely when subjects receive informative feedback, as this helps subjects discern the task structure. Thus, explicit training in the patterns of cue values associated with different classifications could help subjects employ a Bayesian strategy, which relies on combination of all cues. Explicit training about the relationships among cues does help subjects learn to use overall pattern information in other categorization tasks [16]. St. Evans et al. [17] have argued that trial-and-error feedback training likely leads to implicit learning of cue-criterion relationships, whereas explicit training leads to more explicit, conscious knowledge of the cue structure. It is not clear, however, how differences in the nature of the learning process might affect the procedures subjects use later to classify multi-cue objects in a task like threat assessment.

Previous experiments conducted in this laboratory [1] [2] have examined how people perform the task after learning the relationships among cues and threat classification exclusively by trial-and-error (i.e. with accuracy feedback). Although this is an important training mechanism to consider, operators performing naval threat assessment receive extensive explicit instruction in classifying threats from sensor data. The research will determine whether the aspects of the task emphasized in cognitive training determine the decision strategy employed by participants in subsequent classification judgments.

Purpose of study

The proposed experiments investigate how the provision of explicit cue information affects how people learn to use different pieces of information (cues) to make classification judgments. Specifically, the experiment contrasts *pattern-based* training, in which typical patterns of cues are used to distinguish friend from foe, and *cue-based* training, in which detailed diagnostic information about each individual cue is presented. It is hypothesized that pattern-based training is consistent with a Bayesian decision strategy of weighing all cues, whereas cue-based training is consistent with heuristic strategies of selecting the most diagnostic cue as the basis for classification or performing a simple for/against count of evidence. More specifically, pattern-based training may aid subjects in acquiring knowledge

of how all cues contribute to a contact's classification. In contrast, cue-based training allows subjects to clearly recognize the relative predictiveness of all cues, which could suggest to subjects that a one-reason strategy, such as TTB-C, is useful. It is possible, however, that explicit knowledge of all cue validities could also contribute to use of the Bayesian strategy. Although the Bayesian strategy is not modeled in detail here, it is assumed to involve a computational procedure that evaluates contacts in terms of the computed conditional probabilities of friend or foe classifications given the pattern of available cues.

Method

Subjects

Subjects were 24 men and women who were employees of Defence Research and Development Canada - Toronto (DRDC Toronto), students conducting research at DRDC Toronto, or individuals recruited from local universities. All subjects were aged 18 and older, had normal or corrected-to-normal vision, and were unfamiliar with the specific hypotheses and stimulus configurations of the experiments. All received stress pay remuneration for participating. This study, approved by the DRDC Toronto Human Research Ethics Committee, was conducted in conformity with the Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans (HREC).

Materials

The experiment was conducted using the Team and Individual Threat Assessment Network (TITAN) experimental platform, which is a low fidelity threat assessment simulator. The interface presents a radar screen on which “contacts” are presented by asterisk symbols. Each contact corresponds to a single entity around the subject’s “own ship,” which is indicated by a blue circle at the center of the radar screen. Using the computer mouse, the subject can click on (“hook”) a contact, which activates a set of buttons that allow access to information about that contact. This information consists of four characteristics, such as speed, altitude, and so on. The interface can be customized to allow subjects to view all of the contact’s characteristics simultaneously or to restrict subjects to viewing one characteristic at a time. By clicking another button, subjects call up a box in which two possible classifications, “friend” and “foe,” are indicated. Radio buttons under each classification allow the subject to indicate a classification judgment. Further windows open to allow subjects to indicate a confidence judgment and receive feedback concerning their classification accuracy. TITAN was run on Pentium PC computers.

Subjects performed two conditions, one involving pattern-based training and the other cue-based training. Consequently, two sets of 100 contacts (50 friend and 50 foe) were created for the training sessions and two sets of 100 contacts (50 friend and 50 foe) were created for the test sessions. The instructions given to subjects are shown in Annex A.

Design

Three variables were manipulated in this experiment. The first, varied within subjects, was the Cue Validity of each cue used to describe contacts in the training stimuli sets. To vary Cue Validity, each possible value of a cue (values 1 and 2) was probabilistically associated to friend and foe classifications such that each cue differed in diagnosticity. Thus, for one cue each possible value was paired with the friend or foe classification 90% of the time, for another cue 80% of the time and so on. Table 1 indicates the proportions of friend and foe contacts possessing each cue value for the four cues in the two training sets.

Table 1: Relative Frequencies of Cue Values for Friend and Foe Contacts

	SET 1							
	Cue 1 (Direction of Origin)		Cue 2 (Signal Strength)		Cue 3 (Intelligence)		Cue 4 (Comm. Mode)	
	Value 1 (B. Lag.)*	Value 2 (Red Sea)	Value 1 (High)	Value 2 (Medium)	Value 1 (Private)	Value 2 (Platform)	Value 1 (1)	Value 2 (3)
Friend	90%	10%	60%	40%	30%	70%	20%	80%
Foe	10%	90%	40%	60%	70%	30%	80%	20%
	SET 2							
	Cue 1 (Manoeuvre Pattern)		Cue 2 (Counter Measures)		Cue 3 (Electronic Warfare)		Cue 4 (Response)	
	Value 1 (Foxtrot)	Value 2 (Delta)	Value 1 (None)	Value 2 (Jamming)	Value 1 (Big B)**	Value 2 (None)	Value 1 (No Resp)	Value 2 (Given)
Friend	90%	10%	60%	40%	30%	70%	20%	80%
Foe	10%	90%	40%	60%	70%	30%	80%	20%

*B Lag. = Blue Lagoon; ** Big B = Big Bulge

The second variable manipulated was the Contact Type in the test stimuli. Each test set was made up of patterns that offered contrasting predictions of the three contending classification strategies discussed previously; namely TTB-C, the Bayesian strategy, and the Weighted and Unweighted Pros Rules. Eight cue patterns, presented in Table 2, were identified for which at least one strategy offered a differing response than that predicted by the other strategies. The patterns listed in Table 2 refer to the sequence of cue values for the four cues used to describe contacts. Thus, the pattern 1,2,1,1 indicates that Cue 1 possessed Value 1, Cue 2 possessed Value 2, and so on. In Set 1, for example, 1,2,1,1 indicates that the values of the cues Direction of Origin, Signal Strength, Intelligence, and Communication Mode had the specific values of Blue Lagoon, Medium, Private, and 1. In the test set, each of the critical patterns was paired an equal number of times with friend and foe contacts.

From these contacts, we created six Contact Types (A, B, C, D, E, and F). The types distinguish the predicted accuracy of the possible strategies for each item type based on their predicted responses. Thus, Type A and B items, which were the same patterns but differing in the threat class to which they were associated, elicit opposing predictions from TTB-C and the Bayesian strategy. Where TTB-C would predict that these patterns indicate a friend, the Bayesian strategy would predict they indicate a foe, and vice versa. Type C and D patterns elicit the same predictions from TTB-C and the Bayesian strategy but force the Unweighted Pros Rule to guess because equal numbers of cues suggest friend and foe classifications. Types E and F contacts distinguish the Weighted and Unweighted forms of the Pros strategies.

Table 2: Predicted Responses to Contact Types by Hypothesized Strategies

Cue Pattern	Predicted Response of Strategy				Contact Types	
	TTB-C	Bayesian	Weighted Pros	Unweighted Pros	Foe	Friend
1,2,2,2	Foe	Friend	Friend	Friend	B	A
2,1,1,1	Friend	Foe	Foe	Foe	A	B
1,1,1,1	Foe	Foe	Guess	Guess	D	C
2,2,2,2	Friend	Friend	Guess	Guess	C	D
1,1,2,2	Friend	Friend	Friend	Guess	F	E
1,2,1,2	Friend	Friend	Friend	Guess	F	E
2,1,2,1	Foe	Foe	Foe	Guess	E	F
2,2,1,1	Foe	Foe	Foe	Guess	E	F

Note: Cue pattern indicates the value (as 1 or 2) for each cue in order of cues listed in Table 1

In the test set, each of the critical patterns was paired an equal number of times with friend and foe contacts. We predicted the levels of accuracy predicted by the hypothesized decision procedures for each Contact Type, shown in Table 3. The final variable, also varied within subject, was the kind of cognitive training provided.

Table 3: Predicted Accuracy Levels by Contact Type

Heuristic	Contact Type*					
	A	B	C	D	E	F
TTB-C	0%	100%	0%	100%	0%	100%
Bayesian	100%	0%	0%	100%	0%	100%
Weighted Pros Rule	100%	0%	Guess	Guess	0%	100%
Unweighted Pros Rule	100%	0%	Guess	Guess	Guess	Guess

The third variable was the type of Training provided subjects. Two conditions were varied within subjects: Pattern-based Training, in which subjects were taught the most common characteristics of friend and foe contacts, and Cue-based Training, in which subjects received detailed information about the diagnosticity of cues. Because subjects performed the experimental task in both a Cue-based and Pattern-based training condition, two contacts sets were created using different cue labels and cue values but the underlying cue validity

structures of the two sets were the same. The two sets were fully counterbalanced with the two training conditions and the order in which they were performed.

Procedure

The experiment was divided into two sessions for the Cue-based and Pattern-based training conditions, each with a training and test phase. In the training phase, subjects receive cognitive training followed by a trial-and-error practice. In the Pattern-based condition, subjects were given descriptions of a typical friend and typical foe, which indicated the characteristics (cue values) most frequently possessed by friend and foe contacts. In the Cue-based condition, subjects received a table that indicated how each cue was associated to friend and foe contacts. The table contained the information shown in Table 1 and did not specifically highlight the distinction between the prototypical friend and foe contacts.

In the practice session of both conditions, subjects received 100 contacts, of which 50 were friends and 50 foes. The practice set was representative of all possible cue patterns, although some patterns were more likely to occur than others given the structure of cue information. All contacts were presented simultaneously in random positions on the radar screen, although the subject was required to use “zoom in” and “zoom out” buttons to view all of the contacts. Each contact had four cues associated with it, specifying cue values generated according to the probability matrix shown in Table 1. Subjects selected one contact at a time in whatever order they wished and accessed that contact’s cue values. All four values were available on the screen at the same time but the order in which cues were listed was random from contact to contact. Subjects then made a classification judgment, under no time pressure, indicating that the contact is either friend or foe. After this, subjects received accuracy feedback on their classification judgment in the form of a message indicating whether they were correct or incorrect and provision of the correct classification. Subjects received no additional information concerning the predictiveness of cues than that available through the cognitive training. In both the training and test phases, contacts disappeared from the screen immediately after subjects made a judgment to prevent subjects from revisiting the same contact more than once.

Following the training phase, subjects were allowed a short break and then performed the test phase. The test phase followed the same procedure as the training phase with a few important differences. Subjects could no longer access all cue information simultaneously. Instead, each cue was represented by an individual button that subjects pressed to view the value of that cue. The order of the buttons was randomized from contact to contact. Subjects were given no specific instructions concerning how many cues to select; they were told to view whatever cue information they wanted before making their classification judgment. In addition, subjects were presented with 100 contacts (10 each of type A, B, C, D, E, F, and 40 “filler” items) and they did not receive feedback on the accuracy of their judgments. Subjects were given 14 seconds to make their judgments. A clock display counted down the seconds remaining so subjects could monitor how much time remained to make each judgment. If a subject failed to respond within the time limit, the subject’s non-response was counted as an error.

Results

Subjects' performances in the training session were analysed first to determine how well subjects learned to classify contacts. Then, subjects' performance and cue use in the test session were analysed to examine subjects' decision strategies.

Training session

The contacts presented during the training session were divided into four blocks of 25 contacts each, based on the order of presentation (i.e., the first 25 contacts, the next 25, etc.). Accuracy scores (the percentage of contacts correctly classified as friend or foe) were calculated for each block for each subject to create mean accuracy scores, which are shown broken down by Training condition in Figure 3. A two-way, within-subjects Analysis of Variance (ANOVA) revealed significant effects of block [$F(3,69) = 4.33$, $MSe = 0.014$, $p < .01$] and Training condition [$F(1,23) = 7.35$, $MSe = 0.040$, $p < .02$]. There was no significant interaction effect between the two factors [$F(3,69) = 0.096$, $MSe = 0.0005$, $n.s.$].




Figure 3: Classification Accuracy by Block in the Training Session

As can be seen in Figure 3, subjects achieved a high level of accuracy immediately and did not ultimately improve upon it. The main effect of the training block seems to reflect the dip in performance from Block 1 to Block 2. Post-hoc comparisons revealed significant differences between Block 2 of the Pattern-based condition and Block 1 [Tukey HSD $< .05$] and Block 4 [Tukey HSD $< .05$] of the Cue-based condition. Block 2 of the Pattern-based

condition was marginally different from Block 4 of the Cue-based condition [Tukey HSD < .05]. All other comparisons yielded non-significant results. Thus, although performance was somewhat lower in Block 2 of training, this effect likely does not reflect any systematic trend.

Overall, accuracy generally remained constant across Trial Block and subjects' performance overall was near, but not quite, optimal with respect to both the Bayesian and TTB-C strategies. The mean accuracy in the Pattern-based condition of Block 1 [$t(23) = -3.10$, $p < .01$], Block 2 [$t(23) = -4.03$, $p < .01$], Block 3 [$t(23) = -3.19$, $p < .01$], Block 4 [$t(23) = -3.15$, $p < .01$], were all significant lower than the optimum of .92 predicted by TTB-C and, hence, the optimum of .93 predicted by the Bayesian strategy. In the Cue-based condition, only Block 2 differed from .92 [$t(23) = -4.06$, $p < .01$], but Block 1 [$t(23) = -2.41$, $p < .05$], Block 2 [$t(23) = -4.70$, $p < .01$], Block 3 [$t(23) = -2.07$, $p < .06$], Block 4 [$t(23) = -2.16$, $p < .05$], were all at least marginally significantly lower than .93.

Overall, subjects clearly learned to classify contacts to a high degree of accuracy, which indicates learning of individual cue validities. Performance, however, was somewhat better overall when subjects received Cue-based training, in which they received detailed information concerning the predictiveness of each cue.

Classification strategy

The test set was made up entirely of patterns that offered contrasting predictions of the three contending classification strategies (see Table 2). Type A and B items, for example, elicit opposing predictions from TTB-C and the Bayesian strategy. Type C and D patterns elicit the same predictions from TTB-C and the Bayesian strategy but force the Weighted and Unweighted Pros Rule to guess. Type E and F patterns elicit the same predictions from TTB-C, the Bayesian strategy, and the Weighted Pros Rule, but force the Unweighted Pros Rule to guess. Thus, it should be possible to identify post hoc the type of classification strategy employed by subjects by examining their responses to the different patterns.

To assess strategy use by subjects, we examined each subject's proportion of correct responses for each pattern type (A, B, C, D, E, and F) and compared this to the predicted accuracy levels listed in Table 3. A subject was considered to have performed with 100% accuracy for a pattern type if that subject responded correctly to that pattern more often than would be expected by chance, the assumed criterion of which was 5%. For example, the test set contained 10 items with pattern A. Assuming random responding according to a binomial distribution, the probability of selecting seven correct responses is 11.7% whereas the probability of selecting eight correct responses is 4.4%. Consequently, subjects scoring 8, 9, or 10 correct for items with pattern A were scored as showing "100%" performance for those items. Although subjects exhibiting only 8 or 9 correct for these items did not actually score 100%, we concluded that these subjects were consistently applying a strategy that predicted correct response for those items and that any incorrect responses could be attributed to response error. Similarly, we determined that accuracy of 2 items was only 4.4% likely by chance and scored subjects as showing "0%" performance if they correctly responded to 0, 1, or 2 items with pattern A.

Table 4: Response Accuracies by Contact Type and Decision Strategy Classification

Sub.	Cue-Based Training												Pattern-Based Training													
	True Accuracy						Interpolated Accuracy						Decision Strategy	True Accuracy						Interpolated Accuracy						Decision Strategy
	A	B	C	D	E	F	A	B	C	D	E	F		A	B	C	D	E	F	A	B	C	D	E	F	
1	0.5	0.6	0.0	0.8	0.2	1.0	G	G	0	1	0	1	X	0.4	0.9	0.4	0.6	0.1	1.0	G	1	G	G	0	1	X
2	1.0	0.0	0.3	0.6	0.0	1.0	1	0	G	G	0	1	WP	1.0	0.1	0.0	0.9	0.1	0.9	1	0	0	1	0	1	Bayes
3	0.8	0.5	0.7	0.6	0.2	0.8	1	G	G	G	0	1	X	1.0	0.1	0.2	0.5	0.7	0.4	1	0	0	G	G	G	X
4	1.0	0.1	0.7	0.0	0.3	0.8	1	0	G	0	G	1	X	1.1	0.0	0.3	0.4	0.5	0.5	1	0	G	G	G	G	UP
5	0.1	0.9	0.1	0.8	0.1	0.9	0	1	0	1	0	1	TTB-C	0.8	0.2	0.7	0.6	0.6	0.5	1	0	G	G	G	G	UP
6	1.1	0.0	0.1	0.9	0.0	1.0	1	0	0	1	0	1	Bayes	0.9	0.2	0.5	0.6	0.5	0.4	1	0	G	G	G	G	UP
7	1.0	0.3	0.9	0.0	0.1	1.0	1	G	1	0	0	1	X	0.0	0.8	0.0	1.0	0.1	1.0	0	1	0	1	0	1	TTB-C
8	0.2	0.5	0.0	1.0	0.2	1.0	0	G	0	1	0	1	X	0.8	0.1	0.0	1.0	0.0	1.0	1	0	0	1	0	1	Bayes
9	1.0	0.2	0.3	0.2	0.5	0.7	1	0	G	0	G	G	X	1.1	0.1	0.2	0.9	0.0	0.9	1	0	0	1	0	1	Bayes
10	0.7	0.3	0.3	1.0	0.1	0.8	G	G	G	1	0	1	X	0.7	0.1	0.5	0.7	0.4	0.5	G	0	G	G	G	G	X
11	0.6	0.9	0.4	0.6	0.3	0.9	G	1	G	G	G	1	X	0.9	0.1	0.2	0.4	0.4	0.7	1	0	0	G	G	G	X
12	0.4	0.7	0.1	0.9	0.1	1.0	G	G	0	1	0	1	X	1.0	0.1	0.3	0.5	0.2	0.7	1	0	G	G	0	G	X
13	0.5	0.8	0.0	1.0	0.1	0.8	G	1	0	1	0	1	X	1.0	0.0	0.5	0.3	0.5	0.7	1	0	G	G	G	G	UP
14	0.9	0.4	0.6	0.3	0.6	0.4	1	G	G	G	G	G	X	0.8	0.1	0.8	0.7	0.4	0.5	1	0	1	G	G	G	X
15	0.9	0.0	0.1	0.9	0.0	1.0	1	0	0	1	0	1	Bayes	1.0	0.3	0.0	1.0	0.2	1.0	1	G	0	1	0	1	X
16	0.7	0.3	0.6	0.5	0.2	0.9	G	G	G	G	0	1	X	0.9	0.0	0.1	0.8	0.3	0.8	1	0	0	1	G	1	X
17	1.0	0.3	0.5	0.3	0.2	0.8	1	G	G	G	0	1	X	0.9	0.1	0.5	0.5	0.5	0.5	1	0	G	G	G	G	UP
18	1.0	0.1	0.1	0.7	0.1	1.0	1	0	0	G	0	1	X	0.9	0.1	0.7	0.3	0.5	0.9	1	0	G	G	G	1	X
19	0.8	0.3	0.6	0.6	0.2	0.8	1	G	G	G	0	1	X	0.5	0.4	0.3	0.5	0.6	0.5	G	G	G	G	G	G	Guess
20	0.4	0.2	0.5	0.5	0.5	0.4	G	0	G	G	G	G	X	0.8	0.6	0.4	0.6	0.4	0.4	1	G	G	G	G	G	X
21	0.8	0.3	0.6	0.2	0.3	0.9	1	G	G	0	G	1	X	0.9	0.0	0.2	0.5	0.3	0.9	1	0	0	G	G	1	X
22	1.1	0.2	0.4	0.5	0.0	1.0	1	0	G	G	0	1	WP	0.0	0.9	0.0	1.0	0.1	1.0	0	1	0	1	0	1	TTB-C
23	1.0	0.2	0.0	1.0	0.1	1.0	1	0	0	1	0	1	Bayes	0.8	0.3	0.2	1.0	0.3	0.7	1	G	0	1	G	G	X
24	1.0	0.2	0.5	0.3	0.1	1.0	1	0	G	G	0	1	WP	1.1	0.0	0.3	1.0	0.0	0.9	1	0	G	1	0	1	X

Note: G = Guess, X = Unclassifiable, WP = Weighted Pros, UP = Unweighted Pros

Similar criteria were computed for all pattern types. A subject was considered to be responding at “50%”(guessing rate) if they responded accurately to a number of items consistent with a binomial probability greater than .05.

Table 4 presents subjects’ accuracy scores for each contact type and the decision strategy assigned to each subject. Table 5 presents the number of subjects classified as using a given decision strategy.

Table 5: Number of Subjects Classified as Using Hypothesized Decision Strategies

	Decision Heuristic				
Training Cond.	TTB-C	Pros Rule*	Bayesian	Guessing	Unclassified
Cue-based	1	3	3	0	17
Pattern-based	2	5	3	1	13

* In the Cue-based condition, all Pros Rule users exhibited the Weighted version, whereas all Pros Rule users in the Pattern-Based condition exhibited the Unweighted version.

In contrast to previous experiments [2], in which 75-80% of subjects exhibited a classifiable accuracy pattern, only a minority of subjects in the current experiment exhibited such a pattern. It is unclear why this result occurred as the basic design of the current experiment was the same as that of previous experiments. Also contrary to previous results, three subjects in both the Cue- and Pattern-based conditions exhibited responses consistent with a Bayesian strategy. Only three subjects in total seemed to use the TTB-C heuristic, although this heuristic was used by a larger proportion of subjects in previous experiments [2]. Some subjects used the Pros Rule, although the specific form depended on the Training condition. Specifically, three subjects used the Weighted Pros Rule in the Cue-based condition, whereas five subjects used the Unweighted Pros Rule in the Pattern-based condition.

Response time

Response times were measured from the time at which the subject pressed the “Set Leader” button (which called up the Friend/Foe judgment box) to the time he/she indicated a threat classification and pressed the Return key on the computer keyboard. Note that this time includes the time to inspect cues and the time to indicate a response. Although no predictions concerning response times were drawn from the decision strategies under consideration, mean response times were computed for subjects. Generally, subjects took a fair amount of time, on the order of 7 to 8.75 seconds, to indicate their decisions.

A two-factor, within-subjects ANOVA revealed no significant effect of Training condition [$F(1,23) = 0.36$, $MSe = 10.23$, $n.s.$], but significant effect of Item Type [$F(6,138) = 7.90$, $MSe = 0.97$, $p < .01$]. Subjects tended to respond more rapidly to filler items (X items) than the test items (A, B, C, D, E, F). Filler items were ones for which all the decision procedures predicted the same response and they were more common in the training set, which may have made these items easier in some sense. A significant interaction effect of Training condition

and Item types was observed [$F(6,138) = 2.51$, $MSe = 1.01$, $p < .05$]. The difference in average response times for filler and test items was more pronounced in the Cue-based than Pattern-based training condition.

Cue selections

When a subject clicked on a cue button in the test session to inspect the value of that cue, that action was recorded as a “cue selection.” All such cue selections were recorded for every test item to determine which cues subjects inspected and the order in which they inspected them. The hypothesized strategies make different predictions concerning the order in which subjects should select cues. TTB-C predicts that subjects will select the highest validity cue first and rarely, if ever, inspect other cues. The Bayesian and Pros Rule inspect all cues but there is no need to inspect cues in a particular order. Given that the order of cues on the display was random for each contact, these strategies predict essentially random cue selection patterns, although they do not rule out the use of a consistent but idiosyncratic search order by individuals.

Table 6 shows the average number of cue selections made by subjects overall and for the subsets of subjects assigned to the TTB-C, Pros Rule, and Bayesian strategy. Although each contact possessed only four cues, subjects were able to inspect each cue multiple times, so that for any given contact there could be more than four cue selections (i.e. one or more cues was repeated). Subjects inspected slightly more than four cues on average in both Cognitive Training conditions and there was no significant difference between the Cue-based and Pattern-based conditions [$t(23) = 1.18$, *n.s.*].

Table 6: Average Number of Cue Selections by Decision Strategy

		Decision Strategy Assigned		
Training Cond.	Overall	TTB-C	Pros	Bayesian
Cue-based	4.24	1.51 (N=1)	4.87 (N = 3)	4.54 (N = 3)
Pattern-based	4.46	2.98 (N = 2)	4.74 (N = 5)	4.47 (N = 3)

To determine whether the subset of subjects identified as using TTB-C did, in fact, select a single cue as predicted, separate mean number of cue selections was calculated for conditions in which subjects were inferred to have employed TTB-C, the Pros Rule, or the Bayesian strategy on the basis of subjects’ responses to test items. The numbers of subjects exhibiting each strategy were too few to permit statistical analysis, but as can be seen in Table 7 those subjects classified as TTB-C users inspected fewer cues than subjects classified as using a different strategy. The average number of cues inspected by TTB-C users was greater than one, which indicates these subjects did not engage in perfectly frugal search. Subjects using other strategies inspected more than four cues, which is consistent with those strategies but also indicated subjects occasionally inspected the same cue more than once.




Figure 4: Average Number of Cues Selected by Sequence

Subjects were free to inspect cues in any order they chose but all subjects exhibited a tendency to inspect the most valid cue first. To examine this, we identified for each item which cue was inspected first (First Selection). Figure 4 shows the proportion of items for which the subject selected each cue first, where the cues are indicated in descending order of validity. As can be seen, subjects selected the most predictive cue first more frequently than the next most predictive cue. A repeated measures ANOVA revealed a significant effect of cue validity on First Selections [$F(3,69) = 8.72$, $MSe = 0.058$, $p < .01$] but no effect of Cognitive Training [$F(1,23) = 0.21$, $MSe < 0.0001$, *n.s.*]. Thus, in all cases subjects were most likely to select the most predictive cue first, followed by the remaining cues. There were no appreciable differences in the proportion of First Selections that were the second-, third-, and fourth-most valid cues. A significant interaction effect was observed [$F(3,69) = 5.13$, $MSe = 0.017$, $p < .01$], which seems to reflect the finding that the difference in proportions of time the most predictive cue first compared to all others was larger in the Cue-based than Pattern-based condition.

Discussion

The main purpose of this experiment was to explore the impact of explicit training on subjects' probabilistic cue learning and use of decision heuristics in a simulated threat assessment task. The experiment was only partially successful in achieving this purpose.

The results demonstrated that explicit training is suitable for learning to perform the assessment task. Subjects were able to perform to a near optimal level of accuracy immediately without extensive practice with both cue- and pattern-based training. Thus, explicit training seems to be as effective as feedback training in establishing a high level of performance but does not require as much time and effort.

In comparison to previous experiments [1] [2], only a small number of subjects exhibited an interpretable pattern of responding. This result limits the interpretability of the current findings and it was not possible to detect any difference between subjects' behaviour in the cue- and pattern-based conditions. In light of the fact that most subjects did not yield interpretable data, however, it should not be concluded that differences do not necessarily exist.

One important finding was the occurrence of subjects whose data indicated the use of a Bayesian strategy. Three subjects in each of the cue- and pattern-based conditions appeared to use a Bayesian strategy, a result that was not observed for any subject in the previous experiments conducted in this laboratory [1] [2]. Despite the previous findings that many subjects employed TTB-C in this task after feedback training, only three instances of subject data consistent with TTB-C were observed in this experiment. Again, too few subjects were observed to respond in an interpretable way to make a firm conclusion. Nevertheless, it may be that explicit training, whether cue- or pattern-based, better supports a Bayesian strategy. Explicit training has been linked to superior cue integration in similar tasks and cue integration is a requirement for Bayesian responding [15].

Initially, it was hypothesized that pattern-based training might lead to better cue integration and greater use of a Bayesian strategy than cue-based training. The results offer no evidence to support this hypothesis. Indeed, detailed individual cue information of the sort provided in cue-based training may help subjects internalize the overall cue structure as well as, or better than, pattern-based training. Whereas pattern-based training indicated the direction of cue-criterion associations (i.e., which cue value was more associated with friend or foe), the cue-based training also provided indication of the relative predictiveness of each cue, which would be useful in assessing the conditional probabilities of friend or foe given a pattern of cue values. Neither type of training showed any evidence of promoting the use of TTB-C.

As in previous experiments, most subjects inspected all four cues for each contact and even revisited cues on occasion. This was indicated by the finding that, overall, subjects inspected an average of slightly more than four cues per contact. The three subjects categorized as TTB-C users, however, showed more frugal cue selection behaviour. These subjects inspected 1.5 to 3 cues on average per contact. Their search was not as frugal as predicted by

TTB-C, which requires only a single cue per contact, but was more frugal than the cue selection behaviour of subjects categorized as TTB-C users in previous experiments (see [2]).

In conclusion, the results of the current experiment do not offer a clear picture of how explicit training affects probabilistic cue learning and decision strategy use by subjects in a simulated threat assessment task. Nor do the results indicate how different kinds of explicit information might differ in their effects. Further experimentation, however, is warranted as the experiment did yield a novel finding that some subjects can, given explicit training, employ a Bayesian strategy for this task.

References

1. Bryant, D. J. (2003). *Cue validity learning in threat assessment classification judgments*. Defence R&D Canada – Toronto Technical Report (TR 2003-041). Toronto, Ontario: Defence R&D Canada – Toronto, Department of National Defence.
2. Bryant, D. J. (2005). *Heuristic classification in a simulated air threat task*. Defence R&D Canada – Toronto Technical Report (in press). Toronto, Ontario: Defence R&D Canada – Toronto, Department of National Defence.
3. Bryant, D. J., Webb, R. D. G., & McCann, C. (2003). Synthesizing two approaches to decision making in command and control. *Canadian Military Journal*, 4, 29-34.
4. Gigerenzer, G. (1997). Bounded rationality: Models of fast and frugal inference. *Swiss Journal of Economics and Statistics*, 133, 201-218.
5. Klein, G. (1992). *Decision making in complex military environments*. Final Contract Summary Report, prepared for Naval Command, Control and Ocean Surveillance Center. Fairborn, OH: Klein Associates.
6. Gigerenzer, G., Todd, P. M. & The ABC Research Group. (1999). *Simple heuristics that make us smart*. Oxford, UK: Oxford University Press.
7. Tietz, R. (1992). Semi-normative theories based on bounded rationality. *Journal of Economic Psychology*, 13, 297-314.
8. Todd, P. M. (2001). Fast and frugal heuristics for environmentally bound minds. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox*, pp. 51-70. Cambridge, MA: The MIT Press.
9. Todd, P. M., & Gigerenzer, G. (2001). *Precis of Simple heuristics that make us smart*. *Behavioral and Brain Sciences*, 23, 727-780.
10. Broder, A. (2000). Assessing the empirical validity of the "Take the Best" heuristic as a model of human probabilistic inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 1332-1346.
11. Broder, A., & Schiffer, S. (2003). Bayesian strategy assessment in multi-attribute decision making. *Journal of Behavioral Decision Making*, 16, 193-213.
12. Newell, B. R., Weston, N. J., & Shanks, D. R. (2003). Empirical tests of a fast-and-frugal heuristic: Not everyone "takes-the-best". *Organizational Behavior and Human Decision Processes*, 91, 82-96.
13. Newell, B. R., & Shanks, D. R. (2003). Take the best or look at the rest? Factors influencing "one-reason" decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 53-65.

14. Reiskamp, J., & Hoffrage, U. (1999). When do people use simple heuristics, and how can we tell? In G. Gigerenzer, P. M. Todd, & the ABC Research Group (Eds.), *Simple heuristics that make us smart*, pp. 141-167. Oxford, UK: Oxford University Press.
15. Juslin, P., Jones, S., Olsson, H., & Winman, A. (2003). Cue abstraction and exemplar memory in categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 924-941.
16. Lacroix, G. L., Giguère, G., & Larochelle, S. (2005). The origin of exemplar effects in rule-driven categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31, 272-288.
17. St. Evans, J. B. T., Clibbens, J., Cattani, A., Harris, A., & Dennis, I. (2003). Explicit and implicit processes in multicue judgment. *Memory & Cognition*, 31, 608-618.

Annex A

Training for Probabilistic Cue Learning Instructions

Preamble

This experiment investigates how people learn to use different pieces of information (cues) to make judgments about what an object is. In particular, we want to learn how people deal with cues that are not completely reliable. We are using a naval identification and threat assessment computer task to simulate a defence mission aboard a Navy ship. Before we begin the training session, I will describe the task and provide a demonstration of how to operate the simulator.

Instructions

This task may seem complicated at first but we will describe how it works in detail and allow you to practice before beginning the experiment. Please stop me at any time if you have a question.

Let's begin with an overview of the task:

You are playing the role of an operator in the Operations Room of a naval ship, which is displayed as a blue symbol in the center of your radarscope. Surrounding your ship are asterisks called "contacts." These contacts represent traffic detected by your ship's sensors. Your job is to clear all contacts in your ship's vicinity by assessing their threat levels. Until you classify the contact as friend or foe, the threat they pose is unknown. Threat, in this context, means whether the contact is an enemy who could attack the ship. Your goal in performing the task is to learn the characteristics of friend and foe contacts so that you will be able to describe each of these later in the experiment. Each contact's identity can be determined by several information items that you have access to on your workstation. The information items are called "cues" and come from various sensors on the ship. You will have 4 cues to use in making your judgment whether a contact is friend or foe. Friend and foe contacts are different kinds of craft, with different properties that are reflected in different cue values from the sensors. To make your judgments you will first select a contact for classification, review its cues, then indicate whether it is friend or foe.

For you to make accurate judgments, you will have to learn how their different properties, as reflected in cue values, distinguish friend and foe contacts.

The cues will be “binary” variables; that is, each cue can have only one of two possible values. We have prepared a sheet to indicate the possible values for each cue used in the experiment. Please note, this sheet does **not** indicate how the cues are related to threat class. You should not assume that the order in which the cues and their values are printed on the sheet will tell you anything about how the cues relate to the threat class of contacts. The sheet is solely to let you know what cues and cue values to expect to see in the training session – it will be your job to learn **how** to use those cues.

[Depending on training condition:]

Pattern-Based Training Condition

To help you judge whether a contact is a friend or a foe, we have prepared descriptions of the **typical friend** and **typical foe** [*Description Sheet*]. By typical, these descriptions indicate the characteristics that friend and foe contacts most frequently possess. Because contacts vary in their configuration, every friend will not have all the characteristics of the typical friend and every foe will not have all the characteristics of the typical foe. In fact, any given characteristic of a friend or foe could potentially differ from the typical.

The first session will be a practice session in which you will classify contacts and receive feedback on the accuracy of your judgments. It is important to understand that the cues will never be completely reliable – that is, due to variations in the configurations of the contacts, no particular cue will always be 100% associated with either a friend or foe contact. You can use your descriptions of typical friend and foe to help you, but friends and foes in this session will not always have the same characteristics. After making your judgment for a contact, you will receive feedback. This procedure will be repeated for each contact.

Cue-Based Training Condition

To help you judge whether a contact is a friend or a foe, we have prepared a table that shows how each cue is related to friend and foe classifications [*Description Sheet*]. On the sheet provided, we have listed the four cues and their possible values. Next to each characteristic is indicated what percentage of friends possess that

characteristic and what percentage of foes possess it. Note that the cues are independent of one another and the value of one cue in no way affects the value of any other cue.

The first session will be a practice session in which you will classify contacts and receive feedback on the accuracy of your judgments. It is important to understand that the cues will never be completely reliable – that is, due to variations in the configurations of the contacts, no particular cue will always be 100% associated with either a friend or foe contact. You can use your describing the proportions of friends and foes possessing each characteristic to help you. After making your judgment for a contact, you will receive feedback. This procedure will be repeated for each contact.

[Resume For Both Conditions:]

The purpose of the practice session is strictly for you to learn how to identify friend and foe. We will not be looking at your performance in the training practice, so you need not worry about how many contacts you get correct. It is important to bear in mind that, because no cue is 100% associated with either friend or foe, you cannot expect to ever achieve perfect performance. That is, it is not necessarily possible for anyone to be correct all the time. We just want you to learn to judge friend and foe as accurately as possible.

After the practice session, you will complete a test sessions. The test session will be like the training session but with a few important differences. First, you will not receive any feedback about your accuracy. Second, you will make judgments for only 60 contacts. Although, you will select contacts and indicate your classification judgments and confidence ratings in the same way as the training session, during the test session you will have to click on each cue button separately to view the value for that cue. Third, you will have a limit on the amount of time you can take to make your classification judgment. From the moment you hook a contact, you will have 14 seconds to make your decision. If you have not indicated friend or foe after 14 seconds, the contact you have selected will disappear and your response will be scored as an error.

Demonstration

I will now show you the main features of the computer interface and demonstrate how the task is performed.

Main Features

<i>Feature</i>	<i>Description</i>
Pointer	The arrow on the screen is a pointer that allows you to view menu items. You will use the mouse to move the pointer around your screen.
Radar Scope	The black circle within the large gray box is the “radar scope,” which displays your ship and the contacts within its vicinity.
Ownship	Your ship is located in the center of the shaded circle in the radar scope.
Contact	Surrounding your ship are asterisks called “contacts.” A contact is an object detected by your ship’s sensors that appears on your radar scope.
Zoom In & Zoom Out	These buttons allow you to magnify and minimize the range of the radar scope. You can zoom in as close as 1 nm and zoom out as far as 1024 nm. These buttons are used to bring an out-of-range contact into view.
Menus	Next to the radar scope are buttons that permit you to inspect cues for the selected contact.
Information Items	These items are the cues for the selected contact and provide information about the contact. The cues will always appear in a random order for each contact.
Set Contact	This button allows you to view all the cues for the selected contact at the same time.
Set Threat	This button is used when you are ready to submit your threat assessment. It calls up the Classification Menu.
Classification Menu	This menu is used to indicate the selected contact’s threat classification. You make your decision by clicking the appropriate button under the Friend and foe labels.
Confidence Bar	This bar is used to indicate how confident you feel in your classification judgment, on a scale of 0 to 100. You indicate your rating by clicking on the bar and sliding the pointer up or down until the rating shown at the side is the rating you wish to give.
Short and Cumulative	The short average refers to the average error rate of

Averages	the 6 most recent trials. The cumulative average refers to the average error rate across all trials. These averages are provided solely to help you in learning to classify friend and foe.
----------	---

Step-by-Step Instructions

There are several steps involved in clearing a contact from the display:

1. **Selecting a contact:** First, you must choose a contact. There are no rules governing how you make your selection. You can select randomly, or start at the most zoomed in setting, selecting all visible contacts one-by-one, then moving to the next zoomed out setting, and so on.
2. **Viewing cues:** When you select a contact, the “Set Contact” button will become active. Click on this button to call up a window that displays all the cues for the contact. Each cue will be labeled and show the particular value for the selected contact. In the test session, you will have to click on the button for each individual cue to inspect its value.
3. **Threat Assessment:** Once you are finished reviewing the cue information and want to make your threat classification judgment, click on the “Set Threat” button. A small window will appear with two boxes. The red box is labeled “Foe” and the green box is labeled “Friend.” Click on the button under the appropriate label. You can change your choice by clicking on the other button. You must select one or the other choice. During the beginning of the training session, you will have to guess. Complete your decision by clicking the “Done” button.
4. **Confidence Judgment:** After indicating whether the contact is Friend or foe, indicate your level of confidence in your decision. You do this by clicking on the confidence bar then sliding the pointer to the left or right to change the confidence value that is displayed to the side. Your confidence rating should indicate the chance, out of a hundred, that you believe your judgment is correct. Once you have set the rating to the appropriate value, click the “Done” button.
5. **Feedback:** After making your confidence judgment, you will be given feedback consisting of the correct threat classification. A window will appear that indicates what classification you gave the contact, whether that classification is correct or incorrect, and the correct threat classification. This feedback is provided to help you learn how to use the cues provided for contacts to make correct threat classifications. Feedback will be provided only during the initial training session. During the test session, you will receive no feedback.
6. **Accuracy and Speed:** We want you to achieve the best performance possible. During the training phase, you should concentrate on learning how to use the cues. We will not look at your overall accuracy or speed in making judgments for the training phase; the feedback is provided solely to help you. During the test phase, however, you will not receive any feedback and you should try to be as accurate as possible in making your threat classification judgments. Being accurate is more important than being fast.

7. **Test Session:** During the first session, you will learn how to distinguish foe from friend contacts through trial-and-error practice. There will be a subsequent test session in which we will assess your accuracy in making threat classification judgments. The Test Session will have 60 contacts and you will make a classification judgment and confidence rating as in the training session. You will not be able to view all cue information at once; in the test session you must click on each cue individually to see the value for that cue. There will be no feedback provided. There will be a limit on the amount of time you have to make your classification judgment. Once you select a contact, you will have 14 seconds to indicate friend or foe. If you have not entered your judgment after 14 seconds, the contact will be automatically cleared and your response will be scored as an error. The time limit does not apply to the confidence rating.
8. **Describe Friends and Foes:** In the last phase we will ask you to describe the typical friend and typical foe. You will do this by indicating the characteristics (cues) that tended to be associated with friend and foe and estimate the proportion of friends and foes with each of those characteristics. We will give you a sheet listing all the cues to help you in describing friends and foes. In addition, we will ask you several questions about how you performed the task.

List of symbols/abbreviations/acronyms/initialisms

ANOVA	Analysis of Variance
C2	Command and Control
DRDC	Defence Research & Development Canada
MSe	Mean Square Error
PC	Personal Computer
TITAN	Team and Individual Threat Assessment Network
TTB	Take-the-Best
TTB-C	Take-the-Best-for-Classification

Distribution list

Internal Distribution

DRDC Toronto TR 2005-229

1	-	Director General
1	-	Deputy Director General
1	-	Chief Scientist
6	-	Document Library
6	-	David Bryant (author)
1	-	Joseph Baranski
1	-	Carol McCann
1	-	Megan Thompson
1	-	David Smith
1	-	Angela Febbraro
1	-	Donna Pickering
1	-	Anne-Rene Blais
1	-	Don McCreary
1	-	Keith Hendy
1	-	Justin Hollands
1	-	Lochlan Magee
1	-	Stuart Grant
1	-	Peter Kwantes
1	-	Jocelyn Keillor
1	-	Sharon McFadden

External Distribution

DRDC Toronto TR 2005-229

1	-	DRD KIM
1	-	DRD KIM (unbound)
1	-	DRDC
3	-	DGOR
4	-	DGJFD
1	-	ST HP
2	-	D ST CCIS
2	-	D ST L
2	-	D ST A
2	-	D ST M
2	-	D MRS
1	-	D MRS 6
1	-	D MRS 6-2
2	-	DAT
2	-	DAD
2	-	DLSC
2	-	DLPS
2	-	D DCEI
1	-	D LCSPM
1	-	DRDB
1	-	DMSS
1	-	DMSS – 6
1	-	DMSS – 8
1	-	DMPPD
1	-	DMPPD 6
1	-	DMPPD 6 –2
1	-	DSTCCIS
1	-	DSTA
1	-	DSTM
1	-	DSTM – 2

2	-	DLR
2	-	CFEC
2	-	DRDC Atlantic
2	-	DRDC Valcartier
1	-	Canadian Forces Command and Staff College Toronto Attn: Commanding Officer
1	-	CF Maritime Warfare School CFB Halifax Halifax, Nova Scotia Attn: Commanding Officer

UNCLASSIFIED

DOCUMENT CONTROL DATA (Security classification of the title, body of abstract and indexing annotation must be entered when the overall document is classified)		
1. ORIGINATOR (The name and address of the organization preparing the document, Organizations for whom the document was prepared, e.g. Centre sponsoring a contractor's report, or tasking agency, are entered in section 8.) Publishing: DRDC Toronto Performing: DRDC Toronto Monitoring: Contracting:		2. SECURITY CLASSIFICATION (Overall security classification of the document including special warning terms if applicable.) UNCLASSIFIED
3. TITLE (The complete document title as indicated on the title page. Its classification is indicated by the appropriate abbreviation (S, C, R, or U) in parenthesis at the end of the title) Training Effects on the Use of Simple Heuristics in Threat Assessment (U)		
4. AUTHORS (First name, middle initial and last name. If military, show rank, e.g. Maj. John E. Doe.) David J. Bryant		
5. DATE OF PUBLICATION (Month and year of publication of document.) August 2005	6a NO. OF PAGES (Total containing information, including Annexes, Appendices, etc.) 45	6b. NO. OF REFS (Total cited in document.) 17
7. DESCRIPTIVE NOTES (The category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.) Technical Report		
8. SPONSORING ACTIVITY (The names of the department project office or laboratory sponsoring the research and development – include address.) Sponsoring: Tasking:		
9a. PROJECT OR GRANT NO. (If appropriate, the applicable research and development project or grant under which the document was written. Please specify whether project or grant.)		9b. CONTRACT NO. (If appropriate, the applicable number under which the document was written.)
10a. ORIGINATOR'S DOCUMENT NUMBER (The official document number by which the document is identified by the originating activity. This number must be unique to this document) DRDC Toronto TR 2005–229		10b. OTHER DOCUMENT NO(s). (Any other numbers under which may be assigned this document either by the originator or by the sponsor.)
11. DOCUMENT AVAILABILITY (Any limitations on the dissemination of the document, other than those imposed by security classification.) Unlimited distribution		
12. DOCUMENT ANNOUNCEMENT (Any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, when further distribution (beyond the audience specified in (11) is possible, a wider announcement audience may be selected.) Unlimited announcement		

UNCLASSIFIED

UNCLASSIFIED

DOCUMENT CONTROL DATA

(Security classification of the title, body of abstract and indexing annotation must be entered when the overall document is classified)

13. **ABSTRACT** (A brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual.)

(U) Previous research has examined the use of fast and frugal heuristics for threat classification with probabilistic cues. In all previous studies subjects learned to underlying relationships of cues to friend/foe classification through trial-and-error learning. Explicit training, however, is theoretically interesting because it potentially involves the integration of deliberate cognitive learning of the task's underlying cue structure with implicit, procedural learning of specific cue-criterion probabilities. In the experiment, subjects learned to classify contacts in a simulated naval warfare environment and then were tested on sets of contacts that were designed to contrast predictions of several heuristics, including the Take-the-Best-for-Classification (TTB-C) and Pros Rule developed specifically for the threat classification task, as well as a Bayesian strategy based on computation of the conditional probabilities of friend or foe classification given the particular pattern of cues. The interpretability of the results was limited by the large proportion of subjects who exhibited uninterpretable patterns of responding. In contrast to previous experiments, very few subjects employed TTB-C, although more did use the less frugal Pros Rule. The experiment did yield the novel finding that some subjects can, given explicit training, employ a Bayesian strategy for this task.

14. **KEYWORDS, DESCRIPTORS or IDENTIFIERS** (Technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus, e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

(U) Threat assessment, heuristics, training

UNCLASSIFIED

Defence R&D Canada

Canada's Leader in Defence
and National Security
Science and Technology

R & D pour la défense Canada

Chef de file au Canada en matière
de science et de technologie pour
la défense et la sécurité nationale



www.drdc-rddc.gc.ca

